

ABSTRACT OF THE DISCLOSURE

A naming service clustering method which permits naming service load balancing over a set of object references contained within the same cluster such that loads are equitably distributed among servers. Each cluster contains its own unique object binding table
5 which contains object references which each typically represent a single server. Load balancing is performed using a load balance algorithm (e.g., Round Robin). The specific algorithm which is used to perform load balancing is specified upon creation of each naming service cluster. The method provides transparent load balancing among multiple object references in the same cluster. In addition, transparent “failover” among object references in the same
10 cluster is achieved. The method also provides a dynamic addition of new object references into a cluster via binding, along with the dynamic removal of old object references from a cluster via unbinding.